



PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification 6 : G06F 17/30		A1	(11) International Publication Number: WO 00/25237
			(43) International Publication Date: 4 May 2000 (04.05.00)
<p>(21) International Application Number: PCT/US99/25195</p> <p>(22) International Filing Date: 27 October 1999 (27.10.99)</p> <p>(30) Priority Data: 60/105,883 27 October 1998 (27.10.98) US 09/208,118 9 December 1998 (09.12.98) US</p> <p>(71) Applicant: COMPUTER ASSOCIATES THINK, INC. [US/US]; 1 Corporate Associate Plaza, Islandia, NY 11749 (US).</p> <p>(72) Inventors: ISIP, Amando, B., Jr.; 1204 Serenade Lane, Richardson, TX 75081 (US). PARKER, Christopher, F.; 275 Fox Chase Drive North, Oswego, IL 60543 (US). WEAVER, Stephen, J.; 6113 Heritage Lane, Holland, PA 18966 (US).</p> <p>(74) Agent: DEVITO, Victor; Baker & McKenzie, 805 Third Avenue, New York, NY 10022 (US).</p>			<p>(81) Designated States: AE, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, CA, CH, CN, CU, CZ, DE, DK, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IS, JP, KE, KG, KP, KR, KZ, LK, LR, LS, LT, LU, LV, MD, MG, MK, MN, MW, MX, NO, NZ, PL, PT, RO, RU, SD, SE, SG, SI, SK, SL, TJ, TM, TR, TT, UA, UG, UZ, VN, YU, ZW, ARIPO patent (GH, GM, KE, LS, MW, SD, SL, SZ, TZ, UG, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GW, ML, MR, NE, SN, TD, TG).</p> <p>Published With international search report.</p> <p style="text-align: right;">UNDER EXAMIN</p>
<p>(54) Title: METHOD FOR CHECKING TABLESPACES INVOLVED IN REFERENTIAL INTEGRITY</p> <p>(57) Abstract</p> <p>A recover or reload utility determines if a check utility is to be executed. If a check utility is to be executed, the recover or reload records the dependent tablespace names to virtual storage allocated to a job step (1040) instead of to virtual storage allocated to a subtask (1050). After the recover or reload performs its function, if a check utility is next, the database operating system will not freemain the virtual storage (e.g., data area) into which the recover or reload utility has recorded the names of the dependent tables when the subtask ends. Upon termination of the recover or reload utility, the persistent data area is available for the check utility to use.</p>			
<pre> graph TD 1010["1010 parse load control statements"] --> 1020["1020 load named table"] 1020 --> 1030{"1030 Check data identified?"} 1030 -- Yes --> 1040["1040 allocate virtual storage in jobstep"] 1040 --> 1060["1060 Set dependent tables in check pending status"] 1030 -- No --> 1050["1050 allocate virtual storage in subtask"] 1050 --> 1070["1070 parse check data control statement"] 1070 --> 1080["1080 locate data structures naming tablespaces in check pending status"] 1080 --> 1090["1090 Check data"] </pre>			

FOR THE PURPOSES OF INFORMATION ONLY

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

AL	Albania	ES	Spain	LS	Lesotho	SI	Slovenia
AM	Armenia	FI	Finland	LT	Lithuania	SK	Slovakia
AT	Austria	FR	France	LU	Luxembourg	SN	Senegal
AU	Australia	GA	Gabon	LV	Latvia	SZ	Swaziland
AZ	Azerbaijan	GB	United Kingdom	MC	Monaco	TD	Chad
BA	Bosnia and Herzegovina	GE	Georgia	MD	Republic of Moldova	TG	Togo
BB	Barbados	GH	Ghana	MG	Madagascar	TJ	Tajikistan
BE	Belgium	GN	Guinea	MK	The former Yugoslav Republic of Macedonia	TM	Turkmenistan
BF	Burkina Faso	GR	Greece	ML	Mali	TR	Turkey
BG	Bulgaria	HU	Hungary	MN	Mongolia	TT	Trinidad and Tobago
BJ	Benin	IE	Ireland	MR	Mauritania	UA	Ukraine
BR	Brazil	IL	Israel	MW	Malawi	UG	Uganda
BY	Belarus	IS	Iceland	MX	Mexico	US	United States of America
CA	Canada	IT	Italy	NE	Niger	UZ	Uzbekistan
CF	Central African Republic	JP	Japan	NL	Netherlands	VN	Viet Nam
CG	Congo	KE	Kenya	NO	Norway	YU	Yugoslavia
CH	Switzerland	KG	Kyrgyzstan	NZ	New Zealand	ZW	Zimbabwe
CI	Côte d'Ivoire	KP	Democratic People's Republic of Korea	PL	Poland		
CM	Cameroon	KR	Republic of Korea	PT	Portugal		
CN	China	KZ	Kazakhstan	RO	Romania		
CU	Cuba	LC	Saint Lucia	RU	Russian Federation		
CZ	Czech Republic	LI	Liechtenstein	SD	Sudan		
DE	Germany	LK	Sri Lanka	SE	Sweden		
DK	Denmark	LR	Liberia	SG	Singapore		

METHOD FOR CHECKING TABLESPACES
INVOLVED IN REFERENTIAL INTEGRITY

5

FIELD OF THE INVENTION

The present invention relates to database management systems, and particularly to a method for checking tablespaces involved in referential integrity for a database management system.

10

BACKGROUND INFORMATION

A well known database software program is DATABASE 2 (DB2) database software distributed by IBM Corporation. As is known in the art, DB2 operates as a subsystem in a computer system operating under the IBM MVS operating system software. In a DB2 environment, user data resides in DB2 tables which are in tablespaces. A tablespace is, for example, a portion of storage space in a direct access storage device (DASD) such as a disk drive. For exemplary purposes, illustrated below is an order_entry table that would be stored in a tablespace. The order_entry table contains columns: customer_number; product_code; order_number; buyer_name; and ship_to_zip.

Order_Entry Table

25

customer_number	product_code	order_number	buyer_name	ship_to_zip
1111111111	0010	1234500001	John Doe	60606
1111111111	0040	1234500002	Jane Doe	70707
3333333333	0020	1234500003	Bill Smith	90909
2222222222	0030	1234500004	Fred Smith	80808

30

While the above Order_Entry table shows four rows, the table

could have millions of rows for all the orders of a company, for example 4 million rows. The order_entry table also has, for example, three index keys and two foreign keys. An index key is an identifier for a particular row of a table while a
5 foreign key also identifies a row but is also used for referential integrity as described below. For example, in the order_entry table, one index key could be based on Order_Number, another index key based on buyer_name and a third index key based on ship_to_zip. As is known in the art,
10 an index key for a particular table indicates a row identification (RID) and a selected value for the row (e.g., the index key value).

The index key can be used to generate an index for the table
15 which facilitates subsequent searches for particular data in the table. For example, the Order_Entry table would have three indexes (e.g., one for each index key), each index being stored in an indexspace. Similar to a tablespace, an indexspace is, for example, a designated portion of a DASD.

20 Thus, if a user was looking for rows that contain a particular buyer name in the Order_Entry table, the database management system could query the buyer index for the table to identify all occurrences of the buyer name without reading the entire table to locate the rows.

25 DB2 administrators analyze performance characteristics for application programs that access a database table in an attempt to find the optimum index structure for fast access to the database table. The values to be used as an index must be
30 carefully selected because each index results in overhead for the database system. For example, each transaction in a

database table, such as an add or delete, requires that each index for the table also be updated. Thus, it is desirable that the number of indexes for a table be minimized to enhance the performance of application programs. The values to be used as an index for a database table are selected based on, for example, data accessed most frequently by users of the table, generally on-line transaction users. Index keys generally are not based on foreign keys, as foreign keys are used primarily for validation purposes (e.g., constraint enforcement).

As is known in the art, each table in a database may be either a parent table, a child table or both. A child table is related to a parent table via the foreign key value or values contained in columns of the child table. For example, a foreign key value can appear multiple times in a child table (e.g., multiple rows in a child table can have the same foreign key, such as the customer_number and product_code entries in the order_entry table) but each foreign key must be associated with a unique key in a parent table of the child table.

Referential integrity ensures that every foreign key value is valid (e.g., has a corresponding primary key in a parent table). Thus, referential integrity (RI) means that a value in the column of a row in the table is valid when this value also exists in an index of another table. A row should not be in a table if it violates a constraint. As the order_entry table illustrated above has two foreign keys, it has a RI constraint on customer_number and product_code. As is known in the art, when a user of a DB2 database management system

creates a table, the user also defines the constraints for the table (e.g., the user can define the relational integrity criteria). Illustrated below are an exemplary product table and an exemplary customer table (e.g., the parent tables for the foreign keys in the order_entry table).

Product Table

product_code	product_description	retail_price
00010	laptop pc	1000.00
00020	desktop pc	1100.00
00030	office pc	1200.00
00040	lan pc	3500.00
00050	home pc	999.99

The product table show five rows, although the table could have thousands of rows for all of the different products of a company. The product table has, for example, an index based on the column product_code, which values are illustrated in ascending order. The values in the column product_code are each unique since there is only one product code assigned to each product and thus in this table, a product code would not be included more than once. Accordingly, an index for the product table would include the key value (e.g., the stored value in the product_code column) and a RID. The product table index would reside in a DB2 indexspace.

The customer table illustrated below shows four rows, although this table could also have thousands of rows for all of the customers of a company. The customer table has, for example,

an index based on the column `customer_number`, which values are illustrated in ascending order. The values in the column `customer_number` are each unique since there is only one customer number assigned to each customer name and thus a 5 customer number would not be included in this table more than once. Accordingly, an index for the customer table would include the key value (e.g., the value of the column `customer_number`) and a RID. The customer index would also reside in a DB2 indexspace.

10

Customer Table

customer_number	buyer_name	customer_address
1111111111	John Doe	State A
2222222222	Fred Smith	State B
3333333333	Bill Smith	State C
4444444444	Steve Jones	State D

15

As shown by the above tables, all of the rows in the Order_Entry table are valid (e.g., there are no referential integrity constraint violations) because the foreign key values in the column `product_code` of the Order_Entry table also exist in the product table and the values in the column `customer_number` of the Order_Entry table also exist in the customer table.

25

Conventional database management systems, such as DB2, provide the user with the ability to identify specific conditions that a row must meet before it can be added to a table. These conditions are referred to as "constraints" because they 30 constrain the values that a row may include. Constraints include, for example, check constraints and referential

integrity constraints. Check constraints include, for example, qualifying criteria for a particular value, such as a zip code value (e.g., the ship_to_zip value in the Order_Entry table) being in the range of 00000 to 99999. As discussed 5 above, referential integrity constraints ensure that a value in a row of a table is valid when the value also exists in an index of another table.

Constraint enforcement can be performed prior to loading of 10 data into a database table or after data has already been loaded into a database table. When constraint enforcement is performed after loading data into a database table, for example as part of a recovery operation following a hardware of software failure, the constraint enforcement is generally 15 performed by a CHECK utility, such as CHECK DATA by IBM Corp., CHECK PLUS by BMC Software and FASTCHECK by Platinum technology, inc.

Conventional CHECK utilities ensure that data in the table do 20 not violate any constraints that have been established for the table. Constraints can be established at the time the table is generated. For example, constraints can be defined when the table is originally created in the database system and are stored in the DB2 catalog, which can be subsequently queried 25 by a CHECK utility to identify the constraint information. A description of how to read the DB2 catalog is contained in co-pending application Serial No. 09/151,750 entitled DYNAMIC DETERMINATION OF OPTIMAL PROCESS FOR ENFORCING CONSTRAINTS, assigned to the assignee of the present application and which 30 is hereby expressly incorporated by reference.

To perform constraint enforcement, a conventional CHECK utility would, for example, be initialized and identify any applicable constraints for the table to be checked by reading the DB2 catalog, as is known in the art. The CHECK utility would, for example, then read each row of the database table and check for check constraint violations and/or referential integrity constraint violations.

Some DB2 tables have referential integrity constraints. A typical user application system has, for example, three parent tables and thirty one dependent tables. Some customers have as many as twenty one parent tables for one application system. There are various circumstances in which a database table is altered. For example, the table may be recovered from a backup file or reloaded with a new set of data. When a parent table is reloaded or recovered, the dependent tables may be out of synchronization with their respective parent table. For example, a recovered table may have used a backing file that was several days old and thus the recovered table may be out of synch with a dependent table which was in synch with the pre-recovered parent table.

Accordingly, the recover utility or the reload utility must identify all of the dependent tables for each parent tables that has been recovered or reloaded and set these dependent tables to a "check pending" status. Identification of the dependent table can be accomplished, for example, by reading the DB2 catalog as described above. As is known in the art, dependent tables in check pending this status are not accessible to the user programs until a CHECK utility is executed and any referential integrity violations corrected.

When recovering or reloading a tablespace with defined referential integrity constraints, the customer must run the Report Tablespace-Set utility which will generate a listing of all dependent tablespaces that needs to be checked.

5

For a reload operation in conventional database systems, the user needs to code and run a conventional Load utility naming the tables that should be reloaded. The following example is such a control statement to be provided to a conventional Load 10 utility.

```
LOAD RESUME NO REPLACE INTO TABLE CTINLAND.TBHAMEP1  
LOAD RESUME NO REPLACE INTO TABLE CTINLAND.TBHAMEP2  
LOAD RESUME NO REPLACE INTO TABLE CTINLAND.TBHAMEP3
```

15 As the reload operation puts the dependent tables of the reloaded table into a check pending status, the user of the Load utility also needs to code and run a control statement for a conventional CHECK utility naming all the dependent tablespaces. An example of such a control statement to be 20 provided to a conventional CHECK utility is as follows.

CHECK DATA

```
, TABLESPACE JTINLAND.SCHORDCT  
, TABLESPACE JTINLAND.SCHITMDI  
, TABLESPACE JTINLAND.SCHHDINF  
, TABLESPACE JTINLAND.SCHREFL2  
, TABLESPACE JTINLAND.SCHCMAIL  
, TABLESPACE JTINLAND.SCHEBILDG  
, TABLESPACE JTINLAND.SCHFACFW  
, TABLESPACE JTINLAND.SCHEHRAC  
, TABLESPACE JTINLAND.SCHCOATI  
, TABLESPACE JTINLAND.SCHCTGCM  
, TABLESPACE JTINLAND.SCHSSTCK  
, TABLESPACE JTINLAND.SCHSRORD  
, TABLESPACE JTINLAND.SCPEPRCE  
, TABLESPACE JTINLAND.SQOMFGHD  
, TABLESPACE JTINLAND.SQOMFCHI  
, TABLESPACE JTINLAND.SQOMFCMT  
, TABLESPACE JTINLAND.SQOMFOVR  
, TABLESPACE JTINLAND.SQOMFOVU  
, TABLESPACE JTINLAND.SQOMFMPL  
, TABLESPACE JTINLAND.SQOMFOVP  
, TABLESPACE JTINLAND.SQOMFCHE
```

25

30

35

40

5

10

```
, TABLESPACE JTINLAND.SQOMFUNT
, TABLESPACE JTINLAND.SQOMFTRQ
, TABLESPACE JTINLAND.SCHEITMR
, TABLESPACE JTINLAND.SCHORDFW
, TABLESPACE JTINLAND.SCHORHLD
, TABLESPACE JTINLAND.SCHREFRL
, TABLESPACE JTINLAND.SCHRETYP
, TABLESPACE JTINLAND.SCPEPICS
, TABLESPACE JTINLAND.SCHERCYL
, TABLESPACE JTINLAND.SCHCACKW
```

15

For a recovery operation in conventional database systems, the user needs to code and run a conventional recover utility naming the tablespace that should be recovered.

The following example is a control statement for a conventional recover utility.

20

```
RECOVER TABLESPACE JTINLAND.SCHAMEP1
      TABLESPACE JTINLAND.SCHAMEP2
      TABLESPACE JTINLAND.SCHAMEP3
```

25

Further, the user needs to code and run a control statement for a conventional Check utility naming all the dependent tablespaces from, for example, the Report Tablespace-Set utility. An example of a control statement to be provided to a conventional Check utility is as follows.

30

35

40

45

CHECK DATA

```
, TABLESPACE JTINLAND.SCHORDCT
, TABLESPACE JTINLAND.SCHITMDI
, TABLESPACE JTINLAND.SCHHDINF
, TABLESPACE JTINLAND.SCHREFL2
, TABLESPACE JTINLAND.SCHCMAIL
, TABLESPACE JTINLAND.SCHBILDG
, TABLESPACE JTINLAND.SCHFACFW
, TABLESPACE JTINLAND.SCHEHRAC
, TABLESPACE JTINLAND.SCHCOATI
, TABLESPACE JTINLAND.SCHCTGCM
, TABLESPACE JTINLAND.SCHSSTCK
, TABLESPACE JTINLAND.SCHSRORD
, TABLESPACE JTINLAND.SCPEPRCE
, TABLESPACE JTINLAND.SQOMFGHD
, TABLESPACE JTINLAND.SQOMFCHI
, TABLESPACE JTINLAND.SQOMFCMT
, TABLESPACE JTINLAND.SQOMFOVR
, TABLESPACE JTINLAND.SQOMFOUV
```

5 , TABLESPACE JTINLAND.SQOMFMPL
, TABLESPACE JTINLAND.SQOMFOVP
, TABLESPACE JTINLAND.SQOMFCHE
, TABLESPACE JTINLAND.SQOMFUNK
, TABLESPACE JTINLAND.SQOMFTRQ
, TABLESPACE JTINLAND.SCHEITMR
, TABLESPACE JTINLAND.SCHORDFW
, TABLESPACE JTINLAND.SCHORHLD
10 , TABLESPACE JTINLAND.SCHREFRL
, TABLESPACE JTINLAND.SCHRETYP
, TABLESPACE JTINLAND.SCPEPICS
, TABLESPACE JTINLAND.SCHERCYL
, TABLESPACE JTINLAND.SCHCACKW

15 Therefore, recovery of a parent tablespace or reload of a parent table involves substantial efforts by the user of utility to identify all of the dependent tablespaces to be checked.

20 Furthermore, in a mainframe environment running MVS or a similar operating system, when a program executes, that program is considered to be a job for MVS. For example, when a customer executes a utility, the utility is a job step that must be executed. When the recover operation is executed, it 25 is a subtask of the utility job step. Similarly, when a reload operation is executed, it is also a subtask of the utility job step. When a check operation is executed, it is also a subtask of the utility job step. A MVS job may have many tasks. Execution of a subtask requires virtual storage 30 (e.g., memory).

MVS will provide the virtual storage whenever a subtask needs it. For example, when a recover utility or a reload utility identifies the dependent tables associated with the parent 35 table being recovered or reloaded, the utility reads the DB2 catalog and reads the dependent table names into virtual storage allocated for the subtask so that, for example, the

dependent tables can be placed in a check pending status following a recover or reload action. When the subtask terminates, MVS will clean up and any allocated data area (e.g., virtual storage) will be freed by MVS so that another 5 subtask can use the data area. Accordingly, the data contained in the freed storage are (e.g., the names of the dependent tablespaces) is deleted when the data area is freed and thus the deleted information is not available for subsequent use, for example by the check utility that needs to 10 check each dependent table.

SUMMARY OF THE INVENTION

According to an embodiment of the present invention, a recover or reload utility determines if a check utility is to be 15 executed. If a check utility is to be executed, the recover or reload utility can record the dependent tablespace names to persistent virtual storage allocated to a job step instead of to temporary virtual storage allocated to a subtask. After the recover or reload utility performs its function, if a 20 check utility is next, then the database operating system will not free the persistent virtual storage (e.g., data area) into which the recover or reload utility has recorded the names of the dependent tablespaces when the subtask ends.

25 Thus, according to an embodiment of the present invention, when the recover or reload utility identifies the names of the dependent tablespaces, it will record the names into a persistent data area. Upon termination of the recover or reload utility, the persistent data area is available for the 30 check utility to use. Accordingly, the check utility does not have determine the dependent tablespaces to be checked - e.g.,

no additional reading of the DB2 catalog by the check utility is required.

Brief Description of the Drawings

5

Figure 1 illustrates an exemplary flowchart for a method of checking tablespaces according to an embodiment of the present invention for a load utility.

10 Figure 2 illustrates an exemplary flowchart for a method of checking tablespaces according to an embodiment of the present invention for a recover utility.

Detailed Description of the Invention

15

Figure 1 illustrates an exemplary flowchart for a method of checking tablespaces according to an embodiment of the present invention for a load utility. In step 1010, a load control statement is parsed. The control statement may have the form, 20 for example, of

LOAD RESUME NO REPLACE INTO TABLE CTINLAND.TBHAMEP1
LOAD RESUME NO REPLACE INTO TABLE CTINLAND.TBHAMEP2
LOAD RESUME NO REPLACE INTO TABLE CTINLAND.TBHAMEP3
CHECK DATA (JTINLAND.SCHAMEP1, JTINLAND.S CHAMEP2,
25 JTINLAND.SCHAMEP3)

25

As indicated by the above control statement, the load control statement includes a check utility control statement (e.g., "CHECK DATA JTINLAND.SCHAMEP1"). In contrast, conventional 30 load utility control statements do not include a control statement for a check utility. The parsing of the control statement includes, for example, identifying the terms contained in the control statement by the load utility so that

the load utility can perform the load operation on the named table(s), for example in a manner known in the art. The syntax of the load control statement illustrated above is exemplary and any desired format or syntax can be utilized provided, however, that the load utility can process the selected format or syntax. In step 1020, the table named in the control statement is loaded, for example in a manner known in the art.

- 10 In step 1030, it is determined if a check data utility is identified in the control statement. For example, according to an embodiment of the present invention, a check data control statement is included with the load control statement. As a result, when the load utility parses the load utility control statement, the load utility can identify that a check utility will be run after the load utility. If no check data utility is identified in the control statement, then virtual storage (e.g., memory) is allocated for the subtask (e.g., the execution of the load utility). For example, as part of the 20 operation of a conventional load utility, the MVS operating system would allocate temporary virtual storage (e.g., getmain the data structures in the task subpool). As is known in the art, virtual storage allocated for the task subpool is cleared (e.g., freed) when the subtask ends. If a check utility is 25 identified in step 1030, then in step 1040, virtual storage is allocated in the jobstep. For example, the MVS operating system allocates persistent virtual storage (e.g., getmain the data structures in the jobstep subpool).
- 30 Following step 1040 or step 1050, the dependent tables are set to a check pending status, for example in a manner known in

the art. In step 1070, the check data control statement, for example contained in the load control statement, is parsed to start the check data process. In step 1080, the check utility locates the data structures (e.g., virtual storage) naming the 5 tablespaces in check pending status. In step 1090, the check data process is performed, for example in a manner known in the art. Pseudocode for an exemplary embodiment of the process illustrated in Figure 1 is provided below.

```
10      Parse Load control statements
          Load named table
          If check utility is next
              Getmain the data structures in jobstep subpool else
              Getmain the data structures in task subpool endif
15      Set dependent tablespaces in check pending
          If check utility is next
              Parse Check Data control statements
              Check utility locates the data structures naming the
              tablespaces in check pending
20      Perform Check Data process endif
```

Thus, according to an embodiment of the present invention, an interface is provided between the load utility and the check utility to take advantage of the work already performed by the 25 load utility to name the dependent tables needing to be checked by the check utility. As a result, the check utility does not need to read, for example, the DB2 catalog as the load utility has already performed this task and saved the information to persistent virtual storage. Accordingly, a 30 user of the load and check utilities does not need to identify all of the dependent tablespaces to the check utility.

Figure 2 illustrates an exemplary flowchart for a method of checking tablespaces according to an embodiment of the present 35 invention for a recover utility. In step 2010, a recover control statement is parsed. The control statement may have

the form, for example, of

RECOVER TABLESPACE JTINLAND.SCHAMEP1
TABLESPACE JTINLAND.SCHAMEP2
TABLESPACE JTINLAND.SCHAMEP3
5 CHECK DATA JTINLAND.SCHAMEP1,
 JTINLAND.SCHAMEP2, JTINLAND.SCHAMEP3

As indicated by the above control statement, the recover control statement includes a check utility control statement (e.g., "CHECK DATA JTINLAND.SCHAMEP1"). In contrast, conventional recover utility control statements do not include a control statement for a check utility. The parsing of the control statement includes, for example, identifying the terms contained in the control statement by the recover utility so that the recover utility can perform the recover operation on the named tablespace(s), for example in a manner known in the art. The syntax of the recover control statement illustrated above is exemplary and any desired format or syntax can be utilized provided, however, that the recover utility can process the selected format or syntax. In step 2020, the tablespace named in the control statement is recovered, for example in a manner known in the art.

In step 2030, it is determined if a check data utility is identified in the control statement. For example, according to an embodiment of the present invention, a check data control statement is included with the recover control statement. As a result, when the recover utility parses the recover utility control statement, the recover utility can identify that a check utility will be run after the recover utility. If no check data utility is identified in the control statement, then virtual storage (e.g., memory) is allocated for the subtask (e.g., the execution of the recover utility). For example, as part of the operation of a

conventional recover utility, the MVS operating system would allocate temporary virtual storage (e.g., getmain the data structures in the task subpool). As is known in the art, virtual storage allocated for the task subpool is cleared (e.g., freed) when the subtask ends. If a check utility is identified in step 2030, then in step 2040, virtual storage is allocated in the jobstep. For example, the MVS operating system allocates persistent virtual storage (e.g., getmain the data structures in the jobstep subpool).

10

Following step 2040 or step 2050, the dependent tablespaces that were not recovered are identified in step 2060. In step 2070, the dependent tablespaces are set to a check pending status, for example in a manner known in the art. In step 2080, the check data control statement, for example contained in the recover control statement, is parsed to start the check data process. In step 2090, the check utility locates the data structures (e.g., virtual storage) naming the tablespaces in check pending status. In step 2100, the check data process is performed, for example in a manner known in the art.

20
Pseudocode for an exemplary embodiment of the process illustrated in Figure 2 is provided below.

25 Parse Recover control statements
Recover named tablespaces
If Check Data is next
 Getmain the data structures in jobstep subpool else
 Getmain the data structures in task subpool endif
30 Identify dependent tablespaces that were not recovered
Set dependent tablespaces in check pending
If Check Data is next
 Parse Check Data control statements
 Check Data locates the data structures naming the tablespaces
 in check pending
35 Perform Check Data process endif

Thus, according to an embodiment of the present invention, an interface is provided between the recover utility and the check utility to take advantage of the work already performed by the recover utility to name the dependent tablespaces needing to be checked by the check utility. As a result, the check utility does not need to read, for example, the DB2 catalog as the recover utility has already performed this task and saved the information to persistent virtual storage. Accordingly, a user of the check utility does not need to identify all of the dependent tablespaces to the check utility.

WHAT IS CLAIMED IS:

1. A method for checking tablespaces in a database management system, the database management system including a load utility and a check utility, the method comprising the steps of:

5 parsing a load utility control statement;
 determining, as a function of the parsed load
utility control statement, if a check utility is identified;
 if the check utility is identified, allocating a
persistent storage area for use by the load utility;
10 recording, by the load utility, a name of a
tablespace in the persistent storage area; and
 locating the persistent storage area for use by the
check utility.

15

2. A method for checking tablespaces in a database management system, the database management system including a recover utility and a check utility, the method comprising the steps of:

20 parsing a recover utility control statement;
 determining, as a function of the parsed recover
utility control statement, if a check utility is identified;
 if the check utility is identified, allocating a
persistent storage area for use by the recover utility;
25 recording, by the recover utility, a name of a
tablespace in the persistent storage area; and
 locating the persistent storage area for use by the
check utility.




Figure 1




Figure 2

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US99/25195

A. CLASSIFICATION OF SUBJECT MATTER

IPC(6) :G06F 17/30

US CL :707/100, 101, 201, 202

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

U.S. : 707/100, 101, 201, 202

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)


EAST, WEST

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 5,758,357 A (BARRY et al.) 26 MAY 1998, column 6, lines 15-67 and column 7, lines 1-63.	1-2

 Further documents are listed in the continuation of Box C. See patent family annex.

Special categories of cited documents:	*T*	later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention
"A" document defining the general state of the art which is not considered to be of particular relevance		
"E" earlier document published on or after the international filing date	*X*	document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone
"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)	*Y*	document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art
"O" document referring to an oral disclosure, use, exhibition or other means		
"P" document published prior to the international filing date but later than the priority date claimed	*A*	document member of the same patent family

Date of the actual completion of the international search
20 DECEMBER 1999Date of mailing of the international search report
13 JAN 2000Name and mailing address of the ISA/US
Commissioner of Patents and Trademarks
Box PCT
Washington, D.C. 20231
Facsimile No. (703) 305-3230Authorized officer
ANTON FETTING

Telephone No. (703) 305-8449